

EXHIBIT D



California's New AI Laws Focus on Training Data, Content Transparency

October 16, 2024

Artificial intelligence (AI) was a central focus of California lawmakers in September 2024. Home to many of the world's leading AI companies, the state [passed a number of new AI bills](#) at the end of the 2024 legislative session. However, despite signing several of those AI bills into law, California Gov. Gavin Newsom vetoed Senate Bill 1047, one of the country's highest profile efforts to require certain AI developers to implement safety measures.

In this alert, we've provided a preliminary roadmap of how the AI regulatory landscape in California has changed, focusing on two of the new laws, Assembly Bill 2013 and Senate Bill 942, which both take effect January 1, 2026. These two laws are potentially the most impactful and will have the broadest application among technology companies, as both laws require enhanced transparency around generative AI. This alert also touches on Newsom's SB 1047 veto and what might come next.

Key takeaways

- This alert focuses on two of California's newly enacted AI laws, AB 2013 and SB 942, both of which take effect January 1, 2026.
- AB 2013 centers on training data transparency, mandating that developers of generative AI models publicly post on their websites certain required information about the data used to train their models.
- SB 942 requires large developers of generative AI systems to offer AI detection tools and watermarking capabilities to end users in connection with audiovisual content.
- Developers should determine if these laws apply to them (or might apply to them in the near future) and prepare to meet their requirements.
- In some instances, laws with delayed effective dates evolve due to political or policy changes as well as other extenuating factors, but as AB 2013 and SB 942 both have been signed into law, now is the time to prepare.
- California's governor vetoed SB 1047 – a bill intended to prevent catastrophic harms caused by AI and that had garnered significant media attention – because he disagreed with the law's threshold for compliance (targeting the largest AI models), which he determined was not backed by adequate data.

AB 2013: Artificial Intelligence Training Data Transparency

At the core of [AB 2013](#) is a disclosure obligation: From January 1, 2026, developers of AI systems must publicly post on their websites certain information about the data used to train those systems.

Applicability – who needs to comply?

This new law is notable for its broad reach. The requirements of AB 2013 apply to "developers" of any "generative artificial intelligence" system that is made available to Californians for use. A "developer" is any person that "designs, codes, produces or

substantially modifies" an AI system for use by the public. Unlike SB 942, there are no quantitative thresholds that must be met before the law applies. AB 2013 also has a backward-looking component: It applies to every "generative artificial intelligence" system (defined broadly to include any AI systems that can "generate derived synthetic content, such as text, images, video, and audio") available to Californians that was released or substantially modified on or after January 1, 2022, including free and paid services. The term "substantially modifies" means "an update that materially changes an AI system's functionality or performance, including by re-training or fine tuning."

Requirements – what do developers need to do?

Starting on January 1, 2026, developers of generative AI systems covered by this law must publicly post on their websites a "high level summary" of the datasets used to train those systems. These summaries must contain 12 specific pieces of information, including but not limited to the following:

- The sources or owners of the datasets.
- A description of the types of data points within the datasets.
- Whether the datasets are protected by copyright, trademark or patent, or in the public domain.
- Whether the datasets were purchased or licensed.
- A statement about whether the datasets contain personal information.
- If the datasets contain synthetic information.
- If and how the datasets were cleaned or modified by the developer.

Given the breadth and ambiguity around these requirements, they may be subject to legal challenges or modification prior to AB 2013's effective date. A key compliance question will be what level of specificity is required to satisfy the "high level summary" requirement of the law. **Note:** There are exceptions to AB 2013's requirements for AI systems designed for cybersecurity and ensuring physical safety, for the operation of aircraft, and for those that are available only to federal entities for security or defense purposes.

Consequences for noncompliance

AB 2013 does not include a specific enforcement mechanism. The legislative commentary suggests that the law will likely be enforced under California's Unfair Competition Law, which authorizes enforcement by the California attorney general, district attorneys and other government prosecutors. It also provides a private right of action, but only when a plaintiff has been injured and lost money or property as a result of violations of the law.

Implications

AB 2013 follows international regulatory trends requiring greater transparency by developers of generative AI systems, including by mandating disclosures about their training datasets. For example, in the European Union's Artificial Intelligence Act, developers of "general-purpose AI models" are required to maintain a sufficiently detailed summary about the content used for training.

Given the broad applicability of AB 2013, the absence of any quantitative thresholds and the importance of California in the tech industry, this law is significant in terms of the number of developers that will be required to publish reports on their training datasets. This overrides the prevailing industry norm to date of developers maintaining the identity and sources of training datasets as proprietary. Significantly, the level of transparency that the law appears to require may intensify disputes around the use of certain training data.

Before the law takes effect on January 1, 2026, developers should prepare themselves for compliance by doing an internal audit of training data sources used to date, including with respect to any content licensed from third parties or obtained from public sources. Going forward, developers should implement practices and procedures for tracking and approving the use of any training datasets in generative AI systems, including obtaining all information required to be reported under the law. As previously noted, it is possible that these requirements may shift before the law comes into effect, but developers should start planning based on the law's current requirements.

SB 942: California AI Transparency Act

Shifting from training data to the output of AI systems, [SB 942](#) aims to help individuals know when content was created or altered by AI.

Applicability – who needs to comply?

SB 942 is focused on generative AI systems that produce images, video and audio content. Unlike AB 2013, SB 942 has a quantitative threshold. It applies only to “covered providers,” defined as “a person that creates, codes, or otherwise produces a generative artificial intelligence system that has over 1,000,000 monthly visitors or users and is publicly accessible within [California].” Note that the 1 million-user threshold is not limited to users located in California – the threshold is 1 million monthly users, located anywhere, of a system that is available to users in California.

SB 942 does not apply to generative AI systems that:

- Do not meet the quantitative threshold.
- Generate text, code, or other types of AI outputs that are not images, video or audio, meaning that some widely used AI chatbots are excluded.
- Provide exclusively non-user-generated content in video games, television, streaming, movie or interactive experiences.

Finally, SB 942 has one provision that applies to users of a covered AI system if they attempt to circumvent certain requirements, as we explain in the “Downstream contractual obligations” section.

Requirements – what do developers need to do?

Beginning January 1, 2026, covered providers must meet key requirements, which we've outlined below.

‘AI detection tool’

Covered providers must make an “AI detection tool” publicly available to users at no cost, and they are obligated to collect and use feedback on the efficacy of the tool. Among other requirements, the AI detection tool must:

- Allow a user to determine whether an image, video, or audio content was created or altered by the covered provider's AI system.
- Provide the user with metadata or other embedded data showing the provenance of the content (excluding any personal information from that metadata).
- Allow a user to utilize the detection tool by either uploading or linking to content.
- Support access via an API so that users do not need to visit the covered provider's website to use the detection tool. (It is

unclear how users will access this API.)

- Only retain content submitted to the tool to comply with the law's requirements.

Watermarking option

Covered providers must give users the ability to include a manifest disclosure in images, video and audio content they create using the covered provider's AI system. These manifest disclosures must identify the content as AI-generated, be clear, conspicuous and understandable to a reasonable person, and be "extraordinarily difficult to remove."

Latent disclosures

Images, video and audio content created by the covered provider's AI system must include latent disclosures. These disclosures must be detectable by the covered provider's AI detection tool and be "extraordinarily difficult to remove," and they must include, to the extent technically feasible and reasonable, the following information (or provide a link to this information):

- "The name of the covered provider."
- "The name and version number of the [AI] system that created or altered the content."
- "The time and date of the content's creation or alteration."
- "A unique identifier."

Downstream contractual obligations

If a covered provider licenses its AI system to a third party, the covered provider must contractually require the third party licensee to maintain the system's ability to include the latent disclosures discussed above. If the covered provider becomes aware that a third party licensee has modified an AI system to prevent the AI system from producing the latent disclosure, it must revoke the license within 96 hours of discovering the violation. Note that this requires knowledge and does not include affirmative obligations to investigate violations. The law also contains a direct obligation on licensees that have their license revoked not to use the covered AI system after they have been removed.

Consequences for noncompliance

The law sets a penalty of \$5,000 per violation, with each day that a covered provider is in violation deemed a separate violation. It is not clear from the law's text whether, for example, two separate images generated without the required latent disclosures on the same day would be considered a single or separate violations. Enforcement falls to the California attorney general or other city or county-level prosecutors.

If a covered provider revokes the license it granted to a third party licensee for failure to maintain latent disclosures and the licensee continues to use the AI system, the licensee may be responsible not just for breach but subject directly to civil actions brought by government prosecutors, who may pursue injunctive relief and reasonable attorney's fees and costs.

Implications

As AI content becomes increasingly indistinguishable from real images, videos, and recordings, the threat of disinformation and misinformation has grown. In response, state legislators have begun implementing watermarking requirements and other transparency measures designed to make sure that people know when they are encountering AI-generated content or systems. SB

942 continues this trend and pushes it further, requiring very specific technical measures to provide that transparency.

Therefore, the first step in compliance for AI developers is to track two dates: the January 1, 2026 effective date of SB 942, and the date on which an AI system they created will have over 1,000,000 monthly users. Compliance with SB 942 will require covered providers to make significant development changes to both the user interface and back-end operations. The law also requires affirmative changes to contractual provisions. Therefore, AI developers should begin planning these updates as part of their roadmaps to ensure they are ready to comply with SB 942 when they anticipate exceeding the threshold. Finally, AI developers should watch for any legal developments that shift the law's requirements before it becomes effective.

Newsom's veto of SB 1047 – Safe and Secure Innovation for Frontier Artificial Intelligence Models Act

Along with the enactment of new laws came an important veto. Newsom chose to send [SB 1047](#), the highest profile AI bill that reached his desk this year, back to the California State Senate. SB 1047 was different from the other AI bills passed by the California Legislature. It was intended to prevent “critical harms” caused by AI – events like the creation of chemical, biological, or nuclear weapons, or mass casualties or hundreds of millions of dollars in damage caused by attacks on critical infrastructure. The law attempted to strike a difficult balance: to protect against these catastrophic potential risks of AI technology while not stifling innovation. To achieve this balance, SB 1047 applied only to large AI systems that exceeded high quantitative processing and training cost thresholds – hence the reference in the law’s name to frontier models. Developers of models subject to the law would have been required to implement a range of security measures.

[In a letter explaining his veto](#), Newsom took issue with those high quantitative thresholds. He argued that requiring safety measures based solely on the size of a model was not a good approach for preventing harms, and that we should instead evaluate a system’s actual risks by taking into account “whether an AI system is deployed in high-risk environments, involves critical decision-making or the use of sensitive data.” But despite this veto, Newsom continued to advocate for California to lead the nation in AI regulation. While proponents of SB 1047 may not have the votes to override Newsom’s veto with a two-thirds majority in both houses of the Legislature, they are likely to draft alternatives aimed at preventing AI’s worst possible harms.

Cooley resource attorney Ben Gould also contributed to this alert.

This content is provided for general informational purposes only, and your access or use of the content does not create an attorney-client relationship between you or your organization and Cooley LLP, Cooley (UK) LLP, or any other affiliated practice or entity (collectively referred to as “Cooley”). By accessing this content, you agree that the information provided does not constitute legal or other professional advice. This content is not a substitute for obtaining legal advice from a qualified attorney licensed in your jurisdiction and you should not act or refrain from acting based on this content. This content may be changed without notice. It is not guaranteed to be complete, correct or up to date, and it may not reflect the most current legal developments. Prior results do not guarantee a similar outcome. Do not send any confidential information to Cooley, as we do not have any duty to keep any information you provide to us confidential. This content may be considered **Attorney Advertising** and is subject to our [legal notices](#).

Key Contacts

Tracy Rubin Palo Alto	trubin@cooley.com +1 650 843 5568
--------------------------	--------------------------------------

Adam Ruttenberg Washington, DC	aruttenberg@cooley.com +1 202 842 7804
Chris Chynoweth New York	cchynoweth@cooley.com +1 650 843 5372

This information is a general description of the law; it is not intended to provide specific legal advice nor is it intended to create an attorney-client relationship with Cooley LLP. Before taking any action on this information you should seek professional counsel.

Copyright © 2023 Cooley LLP, 3175 Hanover Street, Palo Alto, CA 94304; Cooley (UK) LLP, 22 Bishopsgate, London, UK EC2N 4BQ. Permission is granted to make and redistribute, without charge, copies of this entire document provided that such copies are complete and unaltered and identify Cooley LLP as the author. All other rights reserved.